The Effect of Midlife Wealth Taxes on Life Cycle Savings

Seth J. Chandler

July 17, 2025

Abstract

This paper analyzes the life-cycle consumption, saving, and bequest decisions of an individual facing a pre-announced, one-time wealth tax. Using continuous-time optimal control theory, we solve for the optimal paths of consumption and wealth. A key feature of the model is a discontinuous drop in the state variable (wealth) at the time of the tax, which necessitates the derivation of a jump condition for the co-state variable (the shadow price of wealth). We show that the shadow price of wealth must jump upwards at the tax instant, reflecting the fact that post-tax wealth has a higher marginal value in generating future utility. The model is solved numerically to analyze the behavioral effects of the tax. We find that while a higher tax rate unambiguously reduces the final bequest, its effect on pre-tax savings is ambiguous, depending on the relative strength of a negative substitution effect and a positive income effect. The analysis provides insights into the design of policies involving asset tests, such as financial aid or one-off capital levies. The paper is written with a pedagogic focus, providing detailed economic intuition alongside rigorous mathematical derivations, and includes appendices with functional code for replication.

1 Introduction

The life-cycle hypothesis of consumption and saving, pioneered by Modigliani and Brumberg, provides a foundational framework for understanding how rational, forward-looking individuals plan their economic lives.[1] This theory posits that agents save during their working years and dissave in retirement to smooth consumption, thereby decoupling their spending patterns from their income patterns at any given age.[1, 2] This powerful paradigm has become an essential tool for analyzing how public policies, especially those related to taxation, influence household decisions regarding consumption, labor supply, and wealth accumulation.[2, 3] This paper applies and extends this framework to analyze a specific, yet highly relevant, policy intervention: an anticipated, one-time tax on accumulated wealth.

While stylized, a one-time wealth tax serves as a potent analytical analogue for a range of real-world policies that introduce sharp discontinuities and non-linearities into the incentives for saving. A primary example is asset-based means testing for social benefits. Policies such as college financial aid, Medicaid, and Supplemental Security Income (SSI) often feature asset tests that disqualify individuals whose wealth exceeds a certain threshold. For an agent anticipating the need for such a benefit at a future date (e.g., when a child applies to college), this test functions as a high marginal tax rate on wealth accumulated up to that point. Similarly, governments have historically implemented or considered one-off capital levies, particularly in response to major fiscal needs like wars or economic crises.[4] The anticipation of such a levy, even if its exact form is uncertain, creates behavioral responses that can be understood through the lens of the deterministic model analyzed here.

The central economic question of this paper is: How does a rational agent, who knows the timing and rate of a future wealth tax, adjust their entire life-cycle plan? To answer this, we employ the tools of continuous-time optimal control theory, which provides a rigorous framework for solving dynamic optimization problems.[5, 6, 7] The primary contribution of this work is to provide a complete, pedagogically-oriented solution to a life-cycle problem characterized by a discontinuous state variable. In most economic applications of optimal control, state variables like capital or wealth are assumed to evolve continuously over time.[8, 9] The tax event in our model at time t_1 induces a discrete, downward jump in wealth, violating this standard assumption.

This feature requires an extension of the standard toolkit presented in many economic textbooks.[10, 11] Drawing on more advanced results from the control theory literature [12, 13], we rigorously derive a "jump condition" for the co-state variable—the shadow price of wealth. We demonstrate, both mathematically and intuitively, that this shadow price must jump *upward* at the instant of the tax. This jump precisely reflects the increased marginal value of wealth that has successfully passed through the tax

filter. By providing a detailed economic interpretation of this condition, we aim to make this advanced tool more accessible to a broader audience of economists and legal scholars.

The paper is structured as follows. Section 2 formally presents the model. Section 3 applies Pontryagin's Maximum Principle to derive the complete set of necessary conditions for optimality. Section 4 shows how these conditions are integrated to characterize the solution as the root of a single, nonlinear algebraic equation. Section 5 outlines the numerical methodology used to solve the model. Section 6 investigates the economic effects of the tax. Section 7 concludes. Appendices provide functional code for replication.

2 The Model

We consider the problem of an individual agent who lives for a known, finite period of time, normalized to the interval $[0, t_2]$. The agent has an initial endowment of wealth, w_0, and seeks to choose a path of consumption, $c(t)$, to maximize lifetime utility.

2.1 The Agent's Objective

The agent's total utility is derived from two sources: the flow of consumption over their lifetime and the utility from a bequest of wealth, $w(t_2)$, left at the time of death. The objective function is given by:

$$J = \int_0^{t_2} e^{-\rho t} \frac{c(t)^{1-\gamma_1}}{1-\gamma_1} dt + e^{-\rho t_2} \beta \frac{w(t_2)^{1-\gamma_2}}{1-\gamma_2} \tag{1}$$

The functional form for both consumption and bequest utility is of the Constant Relative Risk Aversion (CRRA) type, a standard and flexible specification in models of intertemporal choice.[2, 14] It is the standard choice in part because it exhibits the desirable property that optimal consumption as a fraction of wealth does not depend on the absolute level of wealth, which greatly simplifies the analysis. The parameters governing preferences are:

- $\rho > 0$: The pure rate of time preference, which captures the agent's impatience.
- $\gamma_1 > 0$: The coefficient of relative risk aversion for consumption. Its inverse, $1/\gamma_1$, is the elasticity of intertemporal substitution (EIS).
- $\beta > 0$: A parameter weighting the utility from the final bequest, often called the bequest motive.
- $\gamma_2 > 0$: The coefficient of relative risk aversion associated with the bequest.

2.2 The Dynamic Budget Constraint (State Equation)

The agent's wealth, $w(t)$, is the single state variable in the model.[15] It accumulates at a constant risk-free interest rate, r, and is depleted by consumption, $c(t)$. The law of motion for wealth is:

$$\dot{w}(t) = rw(t) - c(t) \quad \text{for } t \in [0, t_1) \cup (t_1, t_2] \tag{2}$$

The agent begins with a given initial wealth, $w(0) = w_0$.

2.3 The Anticipated Wealth Tax (State Discontinuity)

The key feature of this model is a one-time, perfectly anticipated wealth tax. At a known time $t_1 \in (0, t_2)$, the government levies a tax at a proportional rate $\tau \in [0, 1)$, creating a discrete downward jump in wealth:

$$w(t_1^+) = (1-\tau)w(t_1^-) \tag{3}$$

where $w(t_1^-)$ and $w(t_1^+)$ denote the wealth immediately before and after the tax, respectively.

3 Optimal Control Solution

To solve this dynamic optimization problem, we employ Pontryagin's Maximum Principle (PMP), which provides a set of necessary conditions that an optimal path must satisfy.[5, 19]

3.1 The Current-Value Hamiltonian

We formulate the problem using the current-value Hamiltonian, H, which can be interpreted as the total instantaneous value generated by the agent's choices:

$$H(w, c, \lambda) = \frac{c^{1-\gamma_1}}{1-\gamma_1} + \lambda(rw - c) \tag{4}$$

Here, $\lambda(t)$ is the co-state variable, representing the shadow price of an additional unit of wealth at time t, measured in units of current-value utility.[15]

3.2 Necessary Conditions from Pontryagin's Maximum Principle

An optimal path must satisfy the following four conditions.

3.2.1 1. Maximization of the Hamiltonian

The first-order condition with respect to consumption c is:

$$\frac{\partial H}{\partial c} = c^{-\gamma_1} - \lambda = 0 \implies c(t) = \lambda(t)^{-1/\gamma_1} \tag{5}$$

This equates the marginal utility of consumption to its opportunity cost, the shadow price of wealth, $\lambda(t)$.

3.2.2 2. The Co-state Equation

The optimal path of the shadow price is governed by the co-state differential equation:

$$\dot{\lambda}(t) = \rho\lambda(t) - \frac{\partial H}{\partial w} = \rho\lambda(t) - r\lambda(t) = (\rho - r)\lambda(t) \tag{6}$$

This is a version of the celebrated Keynes-Ramsey rule, or Euler equation, in continuous time.

3.2.3 3. The Transversality Condition

At the end of the planning horizon, the co-state variable must equal the marginal utility of the bequest:

$$\lambda(t_2) = \frac{\partial}{\partial w(t_2)}\left[\beta\frac{w(t_2)^{1-\gamma_2}}{1-\gamma_2}\right] = \beta w(t_2)^{-\gamma_2} \tag{7}$$

3.2.4 4. The Jump Condition

The discrete jump in the state variable $w(t)$ at time t_1 requires a jump condition for the co-state variable:[12, 13]

$$\lambda(t_1^-) = \lambda(t_1^+)\frac{\partial w(t_1^+)}{\partial w(t_1^-)} = \lambda(t_1^+)(1-\tau)$$

Rearranging provides the intuitive form:

$$\lambda(t_1^+) = \frac{\lambda(t_1^-)}{1-\tau} \tag{8}$$

This condition reveals that the shadow price of wealth must jump *upward* at the instant the tax is levied ($\lambda(t_1^+) > \lambda(t_1^-)$), because post-tax wealth is scarcer and thus marginally more valuable.

4 Deriving the Final Equation for Optimal Bequest

The four necessary conditions from Section 3 define a two-point boundary value problem that can be solved to find the optimal path. The system can be reduced to a single algebraic equation that implicitly defines the entire plan.

4.1 Derivation of the Final Algebraic Equation

The system of differential equations and boundary conditions can be collapsed into a single algebraic equation that links the known initial wealth, w_0, to the unknown optimal terminal wealth, $w(t_2)$. This is achieved by expressing w_0 as the present discounted value (PV) of all lifetime outlays: the consumption stream and the final bequest. This intertemporal budget constraint takes the form:

$$w_0 = A \cdot w(t_2) + B \cdot w(t_2)^p \tag{9}$$

The first term, $A \cdot w(t_2)$, represents the present value of the final bequest. The coefficient A is simply the discount factor from the end of life back to time zero:

$$A = e^{-rt_2} \tag{10}$$

The second term, $B \cdot w(t_2)^p$, represents the present value of the entire lifetime consumption stream. The exponent p arises from the relationship between the utility curvature of consumption and bequests:

$$p = \frac{\gamma_2}{\gamma_1} \tag{11}$$

The coefficient B is a complex constant that consolidates all other parameters of the model. It is derived by integrating the optimal consumption path, $c(s)$, over the agent's lifetime, discounted back to the present value. The full expression for B is:

$$B = \frac{(\beta)^{-1/\gamma_1} e^{\alpha t_2}}{r+\alpha} \left[(1-\tau)^{-1/\gamma_1} (1 - e^{-(r+\alpha)t_1}) + (e^{-(r+\alpha)t_1} - e^{-(r+\alpha)t_2}) \right] \tag{12}$$

where $\alpha = (\rho - r)/\gamma_1$, and we assume $r + \alpha \neq 0$.

4.1.1 On Analytical Solvability

Equation 9 is a nonlinear polynomial in $w(t_2)$. Because the exponent $p = \gamma_2/\gamma_1$ can be any positive real number, there is no general formula to solve for $w(t_2)$ analytically. This is why numerical root-finding methods, as described in Section 5, are necessary.

There is, however, a special case where an analytical solution exists: when the utility curvature for consumption and bequests are identical, i.e., $\gamma_1 = \gamma_2$. In this case, the exponent $p = 1$, and the equation becomes linear in $w(t_2)$:

$$w_0 = A \cdot w(t_2) + B \cdot w(t_2) = (A + B)w(t_2)$$

This can be easily solved for the optimal terminal wealth:

$$w(t_2)^* = \frac{w_0}{A + B}$$

For any other case where $\gamma_1 \neq \gamma_2$, the model must be solved numerically.

5 Numerical Solution and Analysis

Given the intractability of the analytical solution, we turn to numerical methods to solve the model and analyze its properties.

5.1 Numerical Root-Finding

The core computational task is to solve the single nonlinear equation (9) for the optimal terminal wealth, $w(t_2)$. This reduces to a one-dimensional root-finding exercise.[26] We can employ standard, robust algorithms available in scientific computing libraries like SciPy for Python or built-in functions in Wolfram Language. The provided code in the appendices implements this logic by constructing a function that calculates the difference between the initial wealth w_0 and the present value of all outlays for a given $w(t_2)$, and then finds the root of this function. This approach, using numerical integration (`quad`, `NIntegrate`), is highly robust and avoids potential corner-case issues that can arise from using a closed-form solution for the integrals.

5.2 Baseline Parameterization

To conduct a quantitative analysis, we must assign plausible values to the model's parameters. Table 1 presents our baseline parameterization, with values chosen to reflect standard assumptions in the life-cycle savings literature.[2, 16, 18]

Table 1: Baseline Parameterization

Parameter	Description	Value	Justification / Source
w_0	Initial Wealth	1	Normalization
t_1	Time of Tax	20	Represents a mid-life event
t_2	End of Life	40	A 40-year planning horizon
r	Risk-free Rate	0.03	Standard value for long-run real interest rate
ρ	Discount Rate	0.04	Standard assumption: agents are somewhat impatient ($\rho > r$)
γ_1	Consumption RRA	2.0	Widely used value in macro/finance literature
γ_2	Bequest RRA	3.0	Assumes higher risk aversion for bequests
β	Bequest Motive	1.0	A significant, but not dominant, bequest motive
τ	Tax Rate	0.2	A 20% tax rate as a baseline for analysis

6 Economic Interpretation of the Tax

This section presents the central economic analysis of the paper by performing comparative statics on the tax rate, τ.

6.1 Effect on Final Bequest, $w(t_2)$

An increase in the tax rate τ unambiguously reduces the optimal level of the final bequest, $w(t_2)$. The tax reduces the agent's total lifetime resources. Since both consumption and the bequest are normal goods, a reduction in lifetime wealth leads to a reduction in the demand for both.

6.2 Effect on Pre-Tax Savings, $w(t_1^-)$

In contrast, the effect of an increase in the tax rate τ on wealth accumulated just before the tax, $w(t_1^-)$, is theoretically ambiguous. This ambiguity arises from the interplay of two opposing economic forces:

- **The Substitution Effect:** The tax increases the "price" of future wealth, inducing the agent to substitute towards present consumption. This pushes $w(t_1^-)$ **down**.
- **The Income Effect:** The tax makes the agent poorer over their lifetime. To maintain a smoother consumption path, the agent has an incentive to save *more* in the pre-tax period to cushion the blow. This pushes $w(t_1^-)$ **up**.

Which effect dominates depends on the agent's preferences, particularly the desire for consumption smoothing (governed by γ_1).

6.3 Numerical Illustration of the Ambiguity

Table 2 and Figure 1 demonstrate this theoretical ambiguity by solving the model numerically for a range of tax rates under different preference assumptions.

The numerical results confirm the theoretical predictions. In all panels, the final bequest $w(t_2)$ is a strictly decreasing function of the tax rate τ. However, as shown in Figure 1, the behavior of pre-tax wealth $w(t_1^-)$ varies significantly. In the Baseline and Low Smoothing scenarios, the substitution effect dominates, and $w(t_1^-)$ decreases as the tax rate rises. In contrast, in the High Smoothing scenario, the income effect dominates for low tax rates, and agents respond to the impending wealth loss by saving *more* before the tax.

Table 2: Comparative Statics of the Tax Rate (τ)

Optimal Pre-Tax Wealth ($w(t_1^-)$) and Final Bequest ($w(t_2)$) for Various Tax Rates

Panel A: Baseline Parameters ($\gamma_1 = 2.0, \gamma_2 = 3.0$)

Tax Rate (τ)	Pre-Tax Wealth ($w(t_1^-)$)	Final Bequest ($w(t_2)$)
0.0	1.948	1.455
0.1	1.921	1.258
0.2	1.892	1.072
0.4	1.826	0.729
0.6	1.745	0.435

Panel B: High Desire for Smoothing ($\gamma_1 = 4.0, \gamma_2 = 5.0$)

Tax Rate (τ)	Pre-Tax Wealth ($w(t_1^-)$)	Final Bequest ($w(t_2)$)
0.0	1.701	1.099
0.1	1.705	0.963
0.2	1.708	0.832
0.4	1.709	0.587
0.6	1.696	0.366

Panel C: Low Desire for Smoothing ($\gamma_1 = 1.1, \gamma_2 = 1.1$)

Tax Rate (τ)	Pre-Tax Wealth ($w(t_1^-)$)	Final Bequest ($w(t_2)$)
0.0	2.251	2.015
0.1	2.193	1.737
0.2	2.132	1.479
0.4	2.000	1.018
0.6	1.850	0.628

[[Insert diagram here that shows the results of Table 2 graphically. The x-axis should be the Tax Rate (τ) and the y-axis should be the Pre-Tax Wealth ($w(t_1^-)$). The figure should contain three lines, one for each panel (Baseline, High Smoothing, Low Smoothing), clearly labeled. The lines for the Baseline and Low Smoothing cases should be downward sloping, while the line for the High Smoothing case should be slightly upward-sloping for low tax rates before eventually turning down.]]

Figure 1: Effect of Tax Rate (τ) on Pre-Tax Wealth ($w(t_1^-)$) under Different Preference Scenarios. This figure visualizes the data from Table 2, highlighting how the behavioral response of pre-tax savings to the tax depends on the agent's desire for consumption smoothing (γ_1).

7 Conclusion

This paper has developed and solved a continuous-time life-cycle model to analyze an individual's optimal decisions in the presence of an anticipated, one-time wealth tax. By applying optimal control theory to a problem with a discontinuous state variable, we have derived a complete set of necessary conditions that characterize the optimal plan.

The analysis yields several key findings. Methodologically, the discontinuity in wealth induces an upward jump in its shadow price, with the jump condition providing a precise measure of the tax's distortionary impact. Economically, a higher tax rate unambiguously lowers the final bequest. Most significantly, the effect on pre-tax savings is theoretically ambiguous, stemming from a conflict between a substitution effect and an income effect. Our numerical analysis confirms that the net outcome depends critically on the agent's preferences.

These findings have important implications for public policy. Asset-based means tests for social programs can create significant and complex distortions in saving behavior. The fact that the response of pre-tax saving is ambiguous suggests that the impact of such policies is not uniform across the population and highlights the risk of unintended consequences. Future work could extend the framework to incorporate uncertainty or general equilibrium effects.

A Appendix A: Wolfram Language Implementation

```
(*
This code solves the life-cycle model by finding the root of an objective
function. The function ‘getImpliedW0‘ calculates the required initial wealth ‘w0‘
for a given terminal wealth ‘w2‘. The ‘FindRoot‘ command then finds the ‘w2‘
that makes this required ‘w0‘ equal to the actual ‘w0‘ from the parameters.
The use of NIntegrate makes the solution robust to parameter choices.
*)

(* --- Model Parameters --- *)
params = {
   w0 -> 1.0,      (* Initial wealth *)
   t1 -> 20.0,     (* Time of tax *)
   t2 -> 40.0,     (* End of life *)
   r -> 0.03,      (* Interest rate *)
   rho -> 0.04,    (* Discount rate *)
   g1 -> 2.0,      (* Gamma for consumption *)
   g2 -> 3.0,      (* Gamma for bequest *)
   beta -> 1.0,    (* Bequest motive weight *)
   tau -> 0.2      (* Tax rate *)
};

(* --- Function to calculate implied w0 for a given w2 --- *)
getImpliedW0[w2_?NumericQ, p_List] := Module[
  {A, lambda2, lambda1plus, lambda0, c0, c1plus, cPreTax, cPostTax, pv_c1, pv_c2, pv_w2,
     impliedw0},

  A = (rho - r) / g1 /. p;

  (* 1. Work backwards from w2 to find lambda and consumption at key points *)
  lambda2 = beta * w2^(-g2) /. p;
  lambda1plus = lambda2 * Exp[-(rho - r) (t2 - t1)] /. p;
  lambda0 = lambda1plus * (1 - tau) * Exp[-(rho - r) t1] /. p;

  c0 = lambda0^(-1/g1) /. p;
  c1plus = lambda1plus^(-1/g1) /. p;

  (* 2. Define consumption path functions *)
  cPreTax[s_] := c0 * Exp[A * s];
  cPostTax[s_] := c1plus * Exp[A * (s - (t1 /. p))];

  (* 3. Calculate present value of all outlays using robust numerical integration *)
  pv_c1 = NIntegrate[Exp[-(r /. p) * s] * cPreTax[s], {s, 0, t1 /. p}, WorkingPrecision
    -> 20];
  pv_c2 = NIntegrate[Exp[-(r /. p) * s] * cPostTax[s], {s, t1 /. p, t2 /. p},
    WorkingPrecision -> 20];
  pv_w2 = Exp[-(r /. p) * (t2 /. p)] * w2;

  impliedw0 = pv_c1 + pv_c2 + pv_w2;
  Return[impliedw0]
];

(* --- Solve the Model --- *)
(* Find the w2 that makes the implied w0 equal the actual w0 *)
solution = FindRoot[getImpliedW0[w2, params] - (w0 /. params) == 0, {w2, 1.0, 1.1},
    WorkingPrecision -> 20];
w2opt = w2 /. solution;

Print["Optimal Final Bequest w(t2): ", w2opt];

(* --- Recover Pre-Tax Wealth w(t1-) --- *)
getW1Minus[w2opt_?NumericQ, p_List] := Module[
  {A, lambda2, lambda1plus, lambda0, c0, cPreTax, int1},

  A = (rho - r) / g1 /. p;
  lambda2 = beta * w2opt^(-g2) /. p;
  lambda1plus = lambda2 * Exp[-(rho - r) (t2 - t1)] /. p;
  lambda0 = lambda1plus * (1 - tau) * Exp[-(rho - r) t1] /. p;
  c0 = lambda0^(-1/g1) /. p;

  cPreTax[s_] := c0 Exp[A s];

```

```
  int1 = NIntegrate[Exp[(r /. p) * ((t1 /. p) - s)] * cPreTax[s], {s, 0, t1 /. p}];

  (Exp[(r /. p) * (t1 /. p)] * (w0 /. p) - int1)
];

w1minusOpt = getW1Minus[w2opt, params];
Print["Optimal Pre-Tax Wealth w(t1-): ", w1minusOpt];
```

B Appendix B: Python Implementation

```
import numpy as np
from scipy.optimize import root_scalar
from scipy.integrate import quad

#
# This code solves the life-cycle model by finding the root of an objective
# function. ‘present_value_objective‘ computes the difference between the required
# initial wealth for a given terminal wealth ‘w2‘ and the actual initial ‘w0‘.
# ‘root_scalar‘ finds the ‘w2‘ for which this difference is zero. The use of
# ‘quad‘ for numerical integration makes the solution robust.
#

# --- Model Parameters ---
params = {
    'w0': 1.0,      # Initial wealth
    't1': 20.0,     # Time of tax
    't2': 40.0,     # End of life
    'r': 0.03,      # Interest rate
    'rho': 0.04,    # Discount rate
    'g1': 2.0,      # Gamma for consumption
    'g2': 3.0,      # Gamma for bequest
    'beta': 1.0,    # Bequest motive weight
    'tau': 0.2      # Tax rate
}

# --- Function to Solve for w0 given a target w2 ---
def present_value_objective(w2, p):
    """
    Calculates the difference between the required initial wealth for a given
    terminal wealth w2 and the actual initial wealth w0.
    The root of this function is the optimal w2.
    """
    if w2 <= 1e-9: # Ensure w2 is positive
        return p['w0']

    # 1. Use transversality to find lambda(t2) from w2
    lambda2 = p['beta'] * (w2 ** -p['g2'])

    # 2. Solve co-state equation backwards to find lambda at t1+ and 0
    lambda1_plus = lambda2 * np.exp(-(p['rho'] - p['r']) * (p['t2'] - p['t1']))
    lambda0 = lambda1_plus * (1 - p['tau']) * np.exp(-(p['rho'] - p['r']) * p['t1'])

    # 3. Use FOC to find consumption at key points
    c0 = lambda0 ** (-1 / p['g1'])
    c1_plus = lambda1_plus ** (-1 / p['g1'])

    # Define consumption path functions for robust integration
    A = (p['rho'] - p['r']) / p['g1']

    def c_pre_tax(s):
        return c0 * np.exp(A * s)

    def c_post_tax(s):
        return c1_plus * np.exp(A * (s - p['t1']))

    # 4. Calculate the present discounted value of the two consumption streams
    # Integrand is e^(-rs) * c(s)
    pv_c1, err1 = quad(lambda s: np.exp(-p['r'] * s) * c_pre_tax(s), 0, p['t1'])
    pv_c2, err2 = quad(lambda s: np.exp(-p['r'] * s) * c_post_tax(s), p['t1'], p['t2'])

    # 5. Calculate the present discounted value of the bequest
```

```
    pv_w2 = np.exp(-p['r'] * p['t2']) * w2

    # 6. The implied initial wealth is the sum of the present values
    implied_w0 = pv_c1 + pv_c2 + pv_w2

    # Return the difference, which should be zero at the solution
    return implied_w0 - p['w0']

# --- Solve the Model ---
# Find the root of the objective function to get optimal w2
# A safe upper bound for w2 is w0 compounded forward with zero consumption.
lower_bound = 1e-6
upper_bound = params['w0'] * np.exp(params['r'] * params['t2'])

sol = root_scalar(present_value_objective, args=(params,), bracket=[lower_bound,
    upper_bound], method='brenth')
w2_opt = sol.root

print(f"Optimal Final Bequest (w2): {w2_opt:.4f}")

# --- Recover Pre-Tax Wealth w(t1-) ---
def get_w1_minus(w2_optimal, p):
    # This function must re-derive c(0) based on the optimal w2
    lambda2 = p['beta'] * (w2_optimal ** -p['g2'])
    lambda1_plus = lambda2 * np.exp(-(p['rho'] - p['r']) * (p['t2'] - p['t1']))
    lambda0 = lambda1_plus * (1 - p['tau']) * np.exp(-(p['rho'] - p['r']) * p['t1'])
    c0 = lambda0 ** (-1 / p['g1'])

    A = (p['rho'] - p['r']) / p['g1']
    def c_pre_tax(s):
        return c0 * np.exp(A * s)

    # We need the integral of e^(r(t1-s))c(s) from 0 to t1, which is e^(rt1) * PV(pre-
    tax consumption)
    integral_part, err = quad(lambda s: np.exp(-p['r'] * s) * c_pre_tax(s), 0, p['t1'])

    # From the integrated budget constraint: w(t1-) = w0*e^(rt1) - e^(rt1)*PV(pre-tax
    cons)
    w1_minus = np.exp(p['r'] * p['t1']) * p['w0'] - np.exp(p['r'] * p['t1']) *
    integral_part
    return w1_minus

w1_minus_opt = get_w1_minus(w2_opt, params)
print(f"Optimal Pre-Tax Wealth (w1-): {w1_minus_opt:.4f}")
```

References

[1] Arrow, K. J., & Kurz, M. (1970). *Public Investment, the Rate of Return, and Optimal Fiscal Policy*. Johns Hopkins Press.

[2] Attanasio, O. P. (1999). Consumption. In J. B. Taylor & M. Woodford (Eds.), *Handbook of Macroeconomics* (Vol. 1, pp. 741-812). Elsevier.

[3] Brent, R. P. (1973). *Algorithms for Minimization without Derivatives*. Prentice-Hall.

[4] Carroll, C. D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1), 1-55.

[5] De Nardi, M. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies*, 71(3), 743-768.

[6] Dorfman, R. (1969). An economic interpretation of optimal control theory. *The American Economic Review*, 59(5), 817-831.

[7] Gourinchas, P. O., & Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1), 47-89.

[8] Judd, K. L. (1998). *Numerical Methods in Economics*. MIT Press.

[9] Kamien, M. I., & Schwartz, N. L. (1991). *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management* (2nd ed.). North-Holland.

[10] Léonard, D., & Long, N. V. (1992). *Optimal Control Theory and Static Optimization in Economics*. Cambridge University Press.

[11] Liberzon, D. (2011). *Calculus of Variations and Optimal Control Theory: A Concise Introduction*. Princeton University Press.

[12] Modigliani, F., & Brumberg, R. (1954). Utility analysis and the consumption function: An interpretation of cross-section data. In K. Kurihara (Ed.), *Post-Keynesian Economics*. Rutgers University Press.

[13] Pontryagin, L. S., Boltyanskii, V. G., Gamkrelidze, R. V., & Mishchenko, E. F. (1962). *The Mathematical Theory of Optimal Processes*. Wiley-Interscience.

[14] Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (2007). *Numerical Recipes: The Art of Scientific Computing* (3rd ed.). Cambridge University Press.

[15] Virtanen, P., Gommers, R., Oliphant, T. E., et al. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17, 261–272.

[16] Wolfram Research, Inc. (2024). FindRoot. *Wolfram Language & System Documentation Center*. Retrieved from https://reference.wolfram.com/language/ref/FindRoot.html